



RESEARCH ARTICLE

10.1029/2025EA004456

Special Collection:

The Lunar Trailblazer Mission
Collection: Mission, Instruments,
Data Analysis Plan

Key Points:

- HVM³ has completed assembly, testing, calibration, and integration on the Lunar Trailblazer orbiter
- HVM³ radiometric sensitivity is sufficient to confidently discriminate water ice and hydroxyl absorptions on the lunar surface
- HVM³ performance is sufficient to measure 1% absorptions in lunar permanently shadowed regions with terrain-scattered irradiance of 1 W m⁻²

Correspondence to:

D. R. Thompson,
david.r.thompson@jpl.nasa.gov

Citation:

Thompson, D. R., Ehlmann, B. L., Green, R. O., Allen, G. D., Bender, H., Copley-Woods, D., et al. (2026). Calibration and performance of the high resolution volatiles and minerals Moon Mapper (HVM³) on Lunar Trailblazer. *Earth and Space Science*, 13, e2025EA004456. <https://doi.org/10.1029/2025EA004456>

Received 28 APR 2025

Accepted 29 AUG 2025

Author Contributions:














Conceptualization: David R. Thompson, Bethany L. Ehlmann, Robert O. Green, Pantazis Mouroulis, Neil Bowles, Martha House, Rachel L. Klima, Carle Pieters

Data curation: David R. Thompson

Formal analysis: David R. Thompson, Holly Bender, Djuna Copley-Woods, Ian McKinley, Pantazis Mouroulis, Michael Sondheim, Vritika Singh, Christopher Smith, Peter Sullivan

© 2026. Jet Propulsion Laboratory, California Institute of Technology. Government sponsorship acknowledged. Earth and Space Science published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Calibration and Performance of the High Resolution Volatiles and Minerals Moon Mapper (HVM³) on Lunar Trailblazer

David R. Thompson¹ , Bethany L. Ehlmann² , Robert O. Green¹ , Gregory D. Allen¹, Holly Bender¹, Djuna Copley-Woods¹, Michael Eastwood¹, Mark Helmlinger³, Christopher Hummel¹, Jared Keller¹, Andrew Klesh¹, Ian McKinley¹, Bradley D. Moore¹, Pantazis Mouroulis¹, Shriya Nadgauda¹, Michael Sondheim¹, Jose Rodriguez¹, Charles Sarture¹, Calina Seybold¹, Vritika Singh¹, Christopher Smith⁴, Peter Sullivan¹, Quentin Vinckier¹, Walton Williamson¹, Shannon Kian G. Zareh¹, Neil Bowles⁵, Angela M. Dapremont⁶ , Kerri L. Donaldson Hanna⁷ , Christopher S. Edwards⁸ , Emily Felder⁹ , Elise Furlan² , Garni Gharibian^{5,10}, Christopher Haberle⁸, Martha House⁹, Rachel L. Klima⁶ , Jasper Miura² , Carle Pieters¹¹ , Elena Scire² , and Kierra Wilk¹¹ 

¹Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, ²California Institute of Technology, Pasadena, CA, USA, ³National Ecological Observatory Network, Boulder, CO, USA, ⁴Sierra Lobo, Inc., Pasadena, CA, USA, ⁵Department of Physics, Oxford University, Oxford, UK, ⁶John Hopkins Applied Physics Laboratory, Laurel, MD, USA, ⁷Department of Physics, University of Central Florida, Orlando, FL, USA, ⁸Department of Astronomy and Planetary Science, Northern Arizona University, Flagstaff, AZ, USA, ⁹Department of Geology, Pasadena City College, Pasadena, CA, USA, ¹⁰University of California Los Angeles, Los Angeles, CA, USA, ¹¹Department of Earth, Environmental, and Planetary Sciences, Brown University, Providence, RI, USA

Abstract This article reports on the initial calibration and performance of the High-resolution Volatiles and Minerals Moon Mapper (HVM³), slated for launch on the National Aeronautics and Space Administration's Lunar Trailblazer mission. HVM³ is an imaging spectrometer measuring from 600 to 3,600 nm with 10-nm spectral sampling and 50–90 m/pixel ground sampling. The mission goal is to understand the form, abundance, and distribution of water across the lunar surface and the lunar water cycle, accomplished by measuring the distinct absorptions of water ice, adsorbed H₂O, and OH/hydroxyl while controlling for thermal effects with HVM³. HVM³ also has the ability to measure mineralogical composition. HVM³ has been assembled, tested and calibrated in preparation for launch and integrated on the Lunar Trailblazer spacecraft. We review the HVM³ design, calibration process, results, and implications for Lunar Trailblazer science goals. We find HVM³'s radiometric sensitivity is sufficient to confidently measure 1% differences in absorption band strengths under direct solar illumination in single pixel data. In addition, HVM³ has the radiometric precision to discriminate different species of volatile absorptions at irradiances of 1 W m⁻², which will enable mapping and discriminating water ice or other volatiles within most of the Moon's Permanently Shadowed Regions using terrain-scattered illumination.

Plain Language Summary We present laboratory tests of a new instrument, the High-resolution Volatiles and Minerals Moon Mapper (HVM³). HVM³ will launch onboard the National Aeronautics and Space Administration's Lunar Trailblazer mission. It is an imaging spectrometer that will measure the spectrum of reflected sunlight in visible and infrared wavelengths for every pixel of an image. Scientists will interpret these spectra with the goal of measuring how water and ice are distributed across the lunar surface. Our laboratory tests indicate that the instrument is sensitive enough to measure small changes in the content of water on the lunar surface. We find the instrument capable of measuring areas of the Moon that are in permanent shadow using light reflected from neighboring surfaces like crater walls.

1. Introduction

NASA's Lunar Trailblazer mission (Ehlmann et al., 2025) aims to characterize the form, abundance, and distribution of water on the lunar surface and the lunar water cycle. It will orbit the Moon with two remote sensing instruments: the Lunar Thermal Mapper (LTM), a multispectral thermal imager to measure surface temperature and composition, and the High-resolution Volatiles and Minerals Moon Mapper (HVM³), an imaging spectrometer with a spectral range from 600 to 3,600 nm capable of detecting volatile absorption features. Figure 1

Investigation: David R. Thompson, Bethany L. Ehlmann, Robert O. Green, Gregory D. Allen, Holly Bender, Djuna Copley-Woods, Michael Eastwood, Mark Helmlinger, Christopher Hummel, Jared Keller, Andrew Klesh, Ian McKinley, Bradley D. Moore, Shriya Nadgauda, Michael Sondheim, Jose Rodriguez, Charles Sarture, Vritika Singh, Christopher Smith, Peter Sullivan, Quentin Vinckier, Walton Williamson, Shannon Kian G. Zareh, Neil Bowles, Angela M. Dapremont, Kerri L. Donaldson Hanna, Christopher S. Edwards, Elise Furlan, Garni Gharibian, Christopher Haberle, Rachel L. Klima, Jasper Miura, Carle Pieters, Elena Scire, Kierra Wilk

Methodology: David R. Thompson, Bethany L. Ehlmann, Robert O. Green, Gregory D. Allen, Holly Bender, Michael Eastwood, Mark Helmlinger, Christopher Hummel, Jared Keller, Andrew Klesh, Ian McKinley, Bradley D. Moore, Pantazis Mouroulis, Shriya Nadgauda, Michael Sondheim, Jose Rodriguez, Charles Sarture, Calina Seybold, Vritika Singh, Peter Sullivan, Quentin Vinckier, Walton Williamson, Shannon Kian G. Zareh, Neil Bowles, Angela M. Dapremont, Kerri L. Donaldson Hanna, Christopher S. Edwards, Emily Felder, Elise Furlan, Garni Gharibian, Christopher Haberle, Rachel L. Klima, Jasper Miura, Carle Pieters, Elena Scire, Kierra Wilk

Resources: Bethany L. Ehlmann

Software: David R. Thompson

Supervision: Bethany L. Ehlmann, Robert O. Green, Andrew Klesh, Michael Sondheim, Jose Rodriguez, Calina Seybold, Vritika Singh, Walton Williamson, Martha House

Validation: David R. Thompson

Visualization: Emily Felder, Jasper Miura

Writing – original draft: David R. Thompson, Bethany L. Ehlmann, Mark Helmlinger, Michael Sondheim, Angela M. Dapremont, Kerri L. Donaldson Hanna, Emily Felder, Martha House, Rachel L. Klima, Jasper Miura, Carle Pieters, Kierra Wilk

(Left) shows a rendering of the spacecraft with HVM³ visible on a side panel. The image at right shows the as-built HVM³ instrument after assembly. At the time of this writing, the instruments have been assembled, tested, and integrated onto the spacecraft in preparation for launch in 2025. We have completed the pre-launch characterization of instrument radiometry and spectral calibration, enabling an assessment of the detection sensitivity of HVM³ and the capabilities of the instrument for measuring water on the lunar surface.

HVM³ measures lunar spectral surface reflectance and detects absorptions in the spectrum caused by the presence of OH and H₂O (Figure 2). HVM³ will acquire spectral radiance maps of features including highland and mare areas with direct solar illumination and permanently shadowed regions (PSRs) with terrain-scattered solar illumination. The science team will then transform these into I/F maps, remove thermal signals and potentially viewing geometry effects, and finally fit the absorption features with library spectra to determine their shape and depth, revealing the form and relative abundance of water. HVM³ is also capable of detecting minerals such as pyroxene, olivine, and plagioclase that have diagnostic features in its spectral range (Dapremont et al., 2025).

HVM³ is an evolution of the Moon Mineralogy Mapper (M³) imaging spectrometer which discovered and mapped water signatures on the sunlit areas of the Moon using absorption features near 3,000 nm (Pieters et al., 2009). M³'s initial observations were subsequently corroborated by other instruments (Chauhan et al., 2021; Clark, 2009; Sunshine et al., 2009). The HVM³ spectrometer is designed to address several key issues about water on the Moon, which the M³ data and subsequent data sets (e.g., Colaprete et al., 2010; Honniball et al., 2021) leave unresolved:

1. Determining whether the observed spectral absorption feature near 3,000 nm is caused by hydroxyl, adsorbed molecular H₂O, or H₂O in ice form. Unlike M³, the HVM³ spectral range extends to 3,600 nm, capturing the complete absorption features of OH hydroxyl, molecular H₂O, and water ice. This enables improved discrimination of the different forms.
2. Determining if and where PSRs contain water ice. Recent work with M³ data suggested water ice in PSRs (Li et al., 2018). Maps of PSR deposits could guide future scientific sampling or resource extraction efforts on the lunar surface. HVM³ aims to map these features with improved detection sensitivity, enabling direct observations of shadowed terrain using terrain-scattered solar illumination at sub-km spatial scales.
3. Understanding the distribution and change of volatile abundance over time across the lunar surface. Deep Impact spectrometer data indicate possible change in volatile content as a function of time of day (Laferriere et al., 2022; Sunshine et al., 2009). The HVM³ observation plan provides repeat visits of selected polar and midlatitude locations at multiple phase angles. Measuring the lunar surface's volatile deposits at multiple times of the solar day will reveal diurnal variability in water signals.
4. Resolving the contribution of thermal emission effects. M³'s limited spectral range makes it difficult to discriminate H₂O and OH absorptions from changes in thermal emission, so estimates of water abundance are sensitive to thermal modeling assumptions (e.g., Bandfield et al., 2018; Li & Milliken, 2017; Wöhler et al., 2017). In contrast, HVM³ captures the entire water absorption feature to better isolate and remove thermal signals. Moreover, Trailblazer's LTM instrument (Bowles et al., 2025) provides simultaneous direct temperature measurements at 6–100 μm.

This article describes the HVM³'s science objectives, design, and performance. Section 2 reviews the mission design, including the instrument, data product formats, and operations. Section 3 describes the instrument's spectral and spatial response functions. Section 4 describes the algorithms for processing HVM³ data and the laboratory measurements used to characterize its radiometric response. Finally, Section 6 analyzes the as-built performance for a variety of lunar science objectives. We close by discussing the implications for HVM³ analysts. All the code used for HVM³ analyses is available at our public repository (Thompson, 2023).

2. HVM³ Instrument Design

2.1. Hardware and Subsystems

HVM³ is an Offner imaging spectrometer based on the Ultra Compact Imaging Spectrometer (UCIS) design (Van Gorp et al., 2014). Figure 3 shows the optical layout. HVM³ uses an off-axis three-mirror anastigmat telescope with one additional fold mirror. All reflective elements are silver-coated aluminum. The telescope has a




Figure 1. Left: Lunar Trailblazer at the Moon, as rendered by Lockheed Martin and appearing previously in (Bender et al., 2022). Right: The HVM³ optical bench assembly, fully assembled and ready for environmental testing. The vaned surface at the top radiates heat from the spectrometer and cryocooler subsystems. The metal box underneath contains the spectrometer optics.

0.77-milliradian instantaneous field of view, providing ground sampling on the surface of approximately 50–90 m/pixel at the nominal orbit altitude of 100 ± 30 km (77 m at 100 km). Spatial sampling varies proportionally to altitude above the lunar surface. Since the Lunar Trailblazer orbital altitude varies, the footprint of an HVM³ pixel changes throughout the mission.

The spectrometer wavelength range has been adjusted from the original UCIS design, covering the 600–3,600 nm interval in order to disambiguate volatile absorption features due to water species and, if present, other ices (e.g., CO₂, NH₃, H₂S). HVM³ provides approximately 10 nm sampling and Spectral Response Functions (SRFs) with a 16 nm Full Width at Half Maximum (FWHM) spectral resolution. A successful laboratory alignment resulted in excellent cross-track uniformity, with “smile” and clocking that is approximately 16% of one pixel (Bender et al., 2022). This equates to just 8% nonuniformity from center to edge. This low level of spatial nonuniformity is important for accurate measurements of the center wavelengths of volatile absorption features, since it permits a single wavelength calibration to be safely used for the full field of view. Table 1 summarizes the performance parameters. For more information on the HVM³ components and optical design, we refer the reader to a comprehensive description by Bender et al. (2022).

A number of other instrument subsystems enable HVM³'s performance. A thermal subsystem holds the detector at cryogenic temperatures during operations (McKinley, Rodriguez, et al., 2024). It uses a cryocooler with a dedicated electronic control system and closed-loop feedback (the cryocooler electronics, or CCE). The cryocooler rejects heat through the peripheral four blades of a built-in radiator, located at the top of the instrument in Figure 1 (Right). In contrast, the spectrometer and telescope are passively cooled and thermally coupled to the central four blades of the radiator. Nine temperature sensors distributed throughout HVM³ record the temperature of: the cryocooler cold tip, used as a proxy for the focal plane array (FPA); the spectrometer; the spectrometer thermal bus; the cooler mount; the telescope; the telescope baffle; the baseplate; the power supply unit (PSU); and the CCE. The temperature of the FPA is anticipated to be constant through flight operations, maintained by the cryocooler, while the temperature of the passively-cooled spectrometer may change with varying amounts of solar radiation incident on HVM³ as a function of sun angle. This results in a time-varying thermal background signal, which is removed from the science data in processing.

The FPA is a mercury cadmium telluride detector by Teledyne Imaging. It has 640 cross-track elements and 320 active spectral channels (Figure 5). An analog readout circuit, the FPIE, digitizes these signals with a 14 bit dynamic range. A custom PSU provides clean power to the FPA. The FPA is covered by an Order Sorting Filter (OSF) that rejects unwanted higher-order dispersions from the grating; this OSF filter contains three zones with different transmittance properties. These zones have seams at approximately 1,180 nm and 2,160 nm, attenuating some of the light in these channels. Several columns at the edges of the FPA are blocked from external illumination. We use these masked columns to estimate the zero point of the FPA, that is the recorded signal in absence of external illumination, which is useful for calibrating the science data.

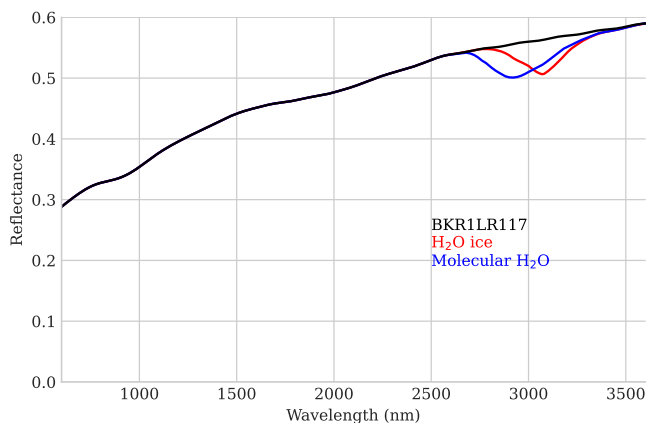


Figure 2. A typical lunar regolith reflectance spectrum from RELAB (Pieters, 1983) with modeled absorptions of molecular H₂O and water ice. At this path length, weaker absorptions at 1,300, 1,500, and 2,000 nm are too small to be visible in this plot.




Figure 3. HVM³ optical design (Bender et al., 2022). Light enters a three mirror telescope where it is imaged onto a 30 μm slit. Light passing through the slit enters an Offner spectrometer where it is dispersed by a concave grating and re-imaged onto a detector array with 30 μm pixel pitch.

2.2. Concept of Operations

Nominal science operation is nadir-observing from orbit with HVM³ and LTM simultaneously and includes the option for a spacecraft roll of up to 8° to enable slight off-nadir pointing. Typically, multiple targeted image acquisitions are planned in each orbit. HVM³ is a pushbroom sensor with a linear field of view, acquiring spatial data line-by-line with full spectral information in each pixel. As the spacecraft orbits, the field of view sweeps out an approximately rectangular area on the lunar surface (Figure 4). The operation begins with an instrument command that is interpreted by the spacecraft to trigger the start of data recording. Instrument digital numbers (DNs) are recorded from the FPIE, and sent to the spacecraft avionics where they are packetized and compressed for transmission to Earth. Upon receipt on Earth by NASA's Deep Space Network, the files are reassembled into data cubes in DN format with two spatial dimensions and a spectral dimension. In addition to science data, Lunar Trailblazer reports HVM³ telemetry data streams used for downstream processing and data calibration.

The instrument has four different operating modes, combining two choices for the image width and two choices for the integration time. The wide image mode reads out all 640 cross-track elements, providing a 24-degree field of view, including masked FPA columns at the periphery of the array. The narrow mode reads out approximately 320 cross-track elements, providing a 12-degree field of view without masked columns. The total size of an HVM³

image is limited by data transfer rates from storage systems in the spacecraft, so the narrow view is used to obtain images with a more balanced aspect ratio of downtrack length to width. The narrow format is also more similar in size to the co-acquired 9-degree LTM footprint. HVM³ can be operated with a 44 ms or 227 ms integration time. The shorter integration time is used for most directly illuminated observations, while the longer integration time is used for low-illumination PSR observations. These longer integrations are combined with five pixel cross track coaddition, coarsening the spatial sampling (by 5 \times 5, e.g. 350 m sampling at 100 km altitude) but improving signal levels for measuring subtle volatile features in shadow. Each integration time requires a separate radiometric calibration.

The instrument has no shutter or cover. In order to estimate the dark state, that is the instrument signal in the absence of photons, HVM³ acquires a “dark image” of an unilluminated scene. During the initial mission phases before Lunar Trailblazer enters science orbit, HVM³ acquires dark images of a patch of deep space that does not contain planets or other bright objects. In its science orbit, HVM³ acquires dark images on the non-sunlit side of the Moon. To accurately remove the instrument dark signal, the dark image integration time must match the integration time of the science data, so in orbits where both PSR and regular mode are used, HVM³ acquires one dark image for each integration time. The dark image consists of many lines - initially 400 - which are reduced into a single frame and then used in subsequent processing of the flight data. As in M³ (Boardman et al., 2011), we schedule dark images in science orbits near the lunar poles to reduce contamination by earthshine.

3. HVM³ Spatial and Spectral Response Functions

During operation, the HVM³ imaging spectrometer acquires sequences of image frames that represent independent reads from the FPA. Each frame is a rectangular array of data recording the light energy incident upon the detector

Table 1
HVM³ Performance Parameters (Bender et al., 2022)

Parameter	Value
Spectral range	600–3,600 nm
Spectral sampling	10 nm
Spectral FWHM	<15 nm
Spatial sampling	50–90 m pixel ⁻¹
Pixel field of view	≤ 0.77 mrad
Spatial along-track response FWHM	0.9 \times sampling interval
Spatial cross-track response FWHM	1.1 \times sampling interval
F-number	3.4
Spatial field of view, full width mode	24°
Spatial field of view, half width mode	12°
Focal length	43 mm
Entrance aperture	12.65 mm
Slit width	30 μm
Detector format	640 \times 480 Pixels
Detector pixel size	30 μm
Raw image format, 24° FOV	640 Columns \times 321 Rows
Radiance image format, 24° FOV	600 Columns \times 308 Rows
Raw image format, 12° FOV	320 Columns \times 321 Rows
Radiance image format, 12° FOV	320 Columns \times 308 Rows
Instrument mass	≤ 14 kg
Telescope temperature, operational	170–196 K
Spectrometer temperature, operational	165–190 K
Detector temperature, operational	118–121 K

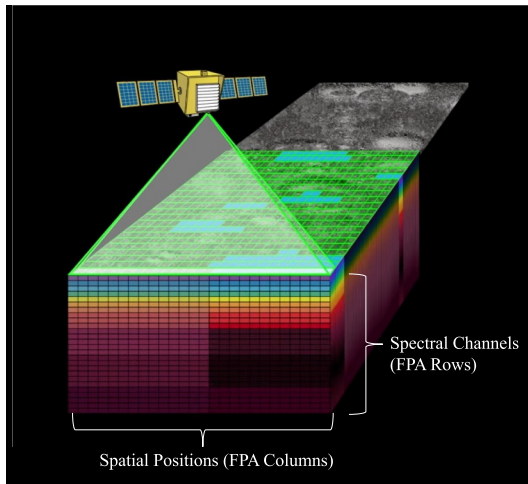


Figure 4. Lunar Trailblazer's linear field of view sweeps across the surface to create a data cube with two spatial dimensions and one spectral dimension. This image is notional; the true detector format contains many more pixels than what is portrayed.

during an integration period. The frame has a spatial dimension and a spectral dimension. As the spacecraft moves over the lunar surface, consecutive frames fill a second spatial dimension to produce a three-dimensional image cube. Each DN value in the cube is related to the intensity of incident light upon the sensor. The calibration process characterizes the spatial and temporal sensitivity of each FPA element by response functions that capture the relative sensitivity distribution of the detector element across each dimension: the SRFs, Cross-track Response Functions (CRFs), and Along-track Response Functions (ARFs).

Each of these response functions is measured in the laboratory using a range of techniques described in Bender et al. (2022). We review the procedures briefly for completeness. The ARFs are formed from the convolution of the telescope Line Spread Function (LSF), representing the blur in the along-track dimension, a rectangular filter, and a second rectangular function representing the motion blur over the finite integration time. ARFs are measured in the laboratory by sweeping a collimated slit across the HVM³ field of view in the downtrack dimension. Tracking the rise and fall of the signal at each cross-track location maps out the response for each downtrack deflection, revealing the ARF. We use Gaussian functions to represent this response. The laboratory measurement also indicates the downtrack deflection of each

pixel's look direction—the angle of the center of its instantaneous field of view relative to the optical axis. The projection of the spectrometer pushbroom slit is not rectilinear with respect to the spacecraft frame; there is a slight curvature which influences the geographic view of each cross-track position.

The CRFs are formed from the end-to-end LSF of the combined spectrometer and telescope system in the cross-track dimension, and the detector pixel response function. The laboratory measurement of CRFs is similar to that of ARFs. We expose the instrument to a spatially narrow collimated beam at multiple field points. We also use Gaussian functions to represent CRFs. However, the properties of spectrometer optics always cause some departure from Gaussianity, a phenomenon often called in-field stray light. The HVM³ instrument prevents this with a new grating groove design that aims to reduce stray spectral and spatial scatter. We assess stray light performance using a bright, collimated slit source that illuminates a single-pixel area of the field. Combining saturated signals (to measure the tails of the response function) with non-saturated signals (to measure the peak) allows characterization of the response across four orders of magnitude.

Figure 6 (Left) shows the result, with a nominal Gaussian response in red. The super-Gaussian tails are approximately two orders of magnitude below the peak response. This behavior is consistent across the spatial field of view and spectral range of the instrument.

The SRFs are most important for spectroscopic interpretation; they describe the sensitivity of each detector element to different wavelengths. Physically, SRFs are formed by the convolution of the rectangular slit function, the spectrometer line spread function, representing the additional blur in the spectral dimension, and the response function of each detector element to photons incident on the array. They are defined by a center wavelength and response function for each channel. In the laboratory, we determine the channel center wavelengths with an integrating sphere illuminated by five lasers with wavelengths spanning the HVM³ spectral range. The Offner spectrometer design has a near-constant dispersion across all wavelengths, so we fit a simple linear function between centroids of the HVM³ responses and the laser wavelengths. Uncertainties are traced to the knowledge of the laser wavelength positions. The laser sphere measurement was repeated at multiple spectrometer temperatures to confirm that the wavelength calibration does not measurably change. Because the lasers illuminate each cross-track element, this measurement is also used to characterize the cross-track uniformity of the

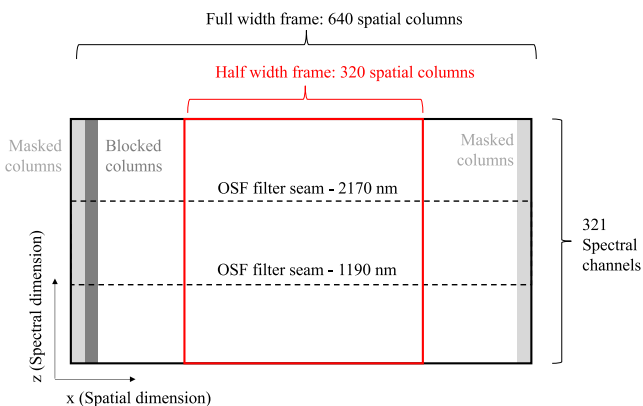


Figure 5. HVM³ focal plane array layout. The sensor can read both half width and full width frames, comprised of 320 and 640 columns respectively. Masked columns at the edge are shielded from incoming photons. “Blocked columns” are blocked at the instrument slit; they are sensitive to light but do not see the external scene.

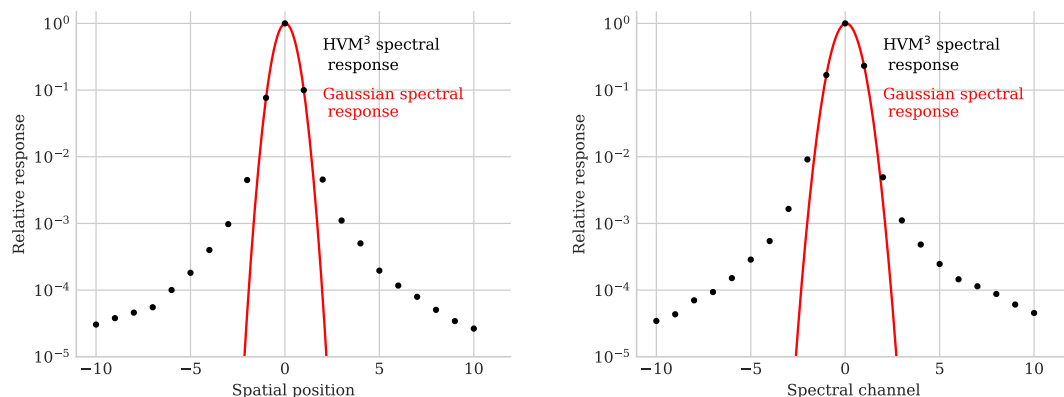


Figure 6. HVM³ spatial (left) and spectral (right) response functions. Black dots represent laboratory measurements of the HVM³ response functions. Red illustrates the nominal gaussian response of the central peak. Spatial and spectral positions are given relative to the element's center location or wavelength, respectively.

spectral calibration. We find that deviations from the nominal center wavelengths are less than 16% of a channel width at all locations on the FPA. For this reason, spectral “smile corrections” of wavelength position should not be needed for routine HVM³ processing.

We determine the HVM³ SRF widths with a monochromator, scanning a collimated monochromatic source over the HVM³ spectral range and recording the output signal. Figure 6 (Right) shows the result. The spectral response is somewhat broader than the CRF relative to the channel spacing, which is an expected attribute of the instrument design. Spectral stray light is approximately two orders of magnitude below the peak response. To obtain the FWHM of the central peak, we fit a Gaussian response function to the rise and fall of the signal in each channel. Figure 7 shows the result of this measurement at three spectrometer field positions. The FWHMs are moderately uniform over the spectral range, and highly uniform across the spatial field of view. FWHMs are slightly larger at the FPA center than at the edges. The periodic structure is most likely an due to variability in the monochromator scan rate, rather than the HVM³ instrument.

4. HVM³ Radiometric Calibration

The HVM³ radiometric calibration relates instrument DNs to radiance at the aperture. In our convention, FPA rows correspond to spectral channels and columns correspond to spatial locations. We use a linear calibration model in which the DNs x , specified at row r and column c , are a function of the radiance at aperture L_{obs} , and gain coefficients k :

$$x(r, c) = k(r, c)L_{obs}(r, c) + x_0(r, c) \quad (1)$$

Here x_0 is the zero point, representing the DN value of each pixel when there are no photons entering the aperture. This value can change over time due to electronic effects, detector dark current, and photons emitted from within the instrument. For convenience we further decompose the gain coefficients into a product of radiometric calibration coefficients (RCCs), denoted a , and a flat field b :

$$k(r, c) = a(r)b(r, c) \quad (2)$$

The RCCs are a function of row only, and represent the sensitivity of the instrument for each wavelength. The flat field is a 2-dimensional array, and represents the spatial variability of the FPA radiometric response. It is convenient to bookkeep flat field and RCCs separately because they generally drift through different processes. The former may shift with electronic changes in individual FPA elements, while the latter generally shifts from

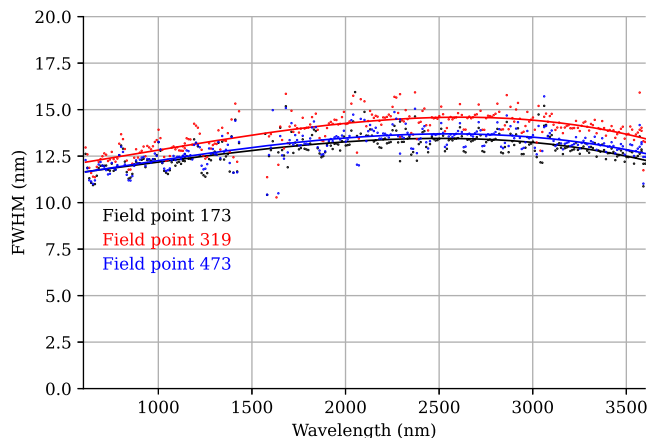


Figure 7. HVM³ spectral response functions measured at three field positions. The full width at half maximum is uniform over the spatial field.

optical degradation or contamination. Separating them with Equation 2 allows the two processes to be tracked and estimated independently. To numerically decouple the flat field from absolute radiometry, each row of b is scaled so that it has an average value of unity within a predefined area of the FPA. This area functions as the reference region where absolute radiometry is measured. For HVM³ this reference area is the central 40 columns, which were the easiest and most reliable to measure in the laboratory calibration environment. Section 4.1 describes our process for estimating the RCCs in detail.

To calibrate an HVM³ frame, we perform the following operations in sequence:

1. Dark frame subtraction—We identify the dark acquisition nearest in time to the science data, and combine it into a single frame using a temporal median. This dark frame is an estimate of the zero point x_0 . We subtract this from each initial DN frame in the data cube.
2. Pedestal shift correction—If this is a wide-format observation with masked columns available, we perform a secondary adjustment to the zero point. For each frame independently, for each row of the FPA, we estimate any residual value in the masked columns using a median of those FPA elements, and then subtract this value uniformly from every pixel in the row. This modifies the original x_0 estimated by the dark frame, capturing any drift since it was acquired. Such drift can be caused by fast-timescale changes to electronic state.
3. Bad pixel interpolation—We identify and replace bad pixels in the FPA. First, we estimate the scene covariance using a robust (i.e., outlier-insensitive) covariance estimator, principal component pursuit by alternating directions (Candès et al., 2011). We use this low rank representation to denoise the spectrum, and then flag as bad any FPA elements that diverge sufficiently far from the denoised version. The principal component procedure also produces a mean and low-rank covariance matrix defining a multivariate Gaussian distribution of spectral variability. We infer appropriate replacements for bad channels in each spectrum independently by calculating the mean of the Gaussian conditioned on the values of the good channels. Each image has its own bad pixel map to indicate which values have been replaced.
4. Thermal correction—Typically, spectrometer self-emission is present in the dark data and fully removed by step 1. However, if the spectrometer temperature varies significantly between dark and illuminated scenes, it is possible that the thermal signal in illuminated data could be higher or lower, creating errors in the zero point estimate. Fortunately, the thermal emission efficiency w_{obs} has a distinctive spectral signature, enabling the magnitude of the residual to be estimated and removed from the data using spectral matching techniques. Section 6 discusses methods for identifying and removing a known spectral perturbation of unknown magnitude using statistics of the background radiance.
5. Flat fielding—We multiply each frame of HVM³ by the flat field frame representing the spatial variability in FPA intensity. We characterized the flat field in the laboratory for the complete field of view in the midwave region above 2,500 nm. However, from 600 to 2,500 nm, stage rotation constraints limited acquisition to the center 50% of the flat field's spatial dimension. We will estimate the remaining portion of the flat field in flight using techniques described in 5.
6. Radiometric calibration—We multiply each spectrum by the RCC vector, translating the data cube into units of $\mu\text{Wnm}^{-1}\text{cm}^{-2}\text{sr}^{-1}$.

After radiometric calibration of the frame is complete, the frame is then clipped to remove any masked columns. The raw data stores longest wavelengths first, but we reverse this in processing so that channels of calibrated radiance data are ordered from shortest wavelength to longest.

A final step of the calibration process, exercised at the option of the science team, aims to correct any pixelwise drift in radiometric gain over time. Instruments like EMIT using similar detectors have measured random gain drifts of approximately 0.1% per month attributed to minor changes in the state of electronic components. An effective solution, outlined in Thompson et al. (2024), fits a local flat field using a large set of recent images. As many as several hundred recent images are combined in a temporal median with high-pass spatial filtering to smooth out scene content and image-based edge detection to remove contours. The remaining scene content averages out to a uniform field over time creating a homogenous stimulus to measure local column-to-column differences in FPA response. This local flat field adjustment addresses any radiometric drift that would appear as striping artifacts. The changes are typically at the sub-percent level, so the impact is mostly cosmetic. Lunar Trailblazer may also periodically collect intentional flat fields using the Moon, as described in Section 5.

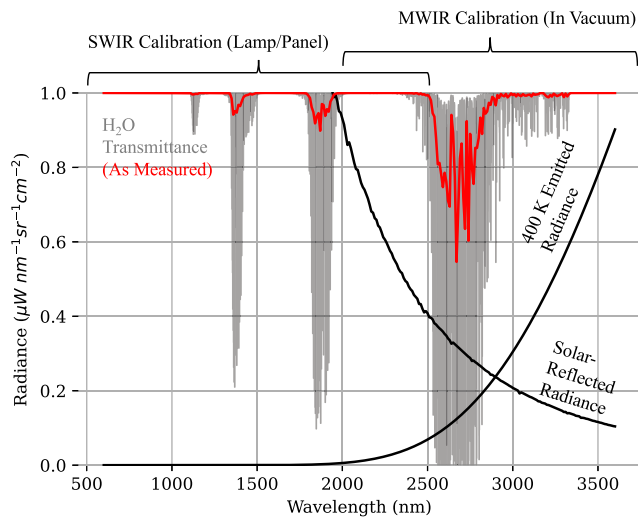


Figure 8. HVM³ uses two calibration methods, a NIST-calibrated lamp for the shortwave range and a temperature-controlled black silicon panel for the midwave range. Gray lines indicate the transmittance of H₂O along the laboratory free air path.

4.1. Measuring Radiometric Calibration Coefficients

The RCCs are a multiplicative factor describing the relationship between the scene signal in instrument DNs, and radiance in $\mu\text{W nm}^{-1} \text{sr}^{-1} \text{cm}^{-2}$. We calculate these values in the laboratory by presenting the instrument with a source of known intensity. Because these data also indicate the optical efficiency, the same laboratory data inform instrument models that can predict the signal to noise ratio (SNR) for other radiances. The uncertainty of these measurements is also recorded for use in downstream analyses.

The HVM³ instrument is calibrated in a vacuum chamber. For most laboratory measurements, it views external calibration equipment through a sapphire window in this chamber. However, the free-air part of this optical path is subject to absorption by atmospheric water vapor in the critical 2,500–3,000 nm range. This absorption is difficult to characterize and would invalidate precise radiometric measurements. Consequently, to ensure that no water vapor signals are present in the radiometric calibration, we calibrate wavelengths longer than 2,500 nm using a separate source in vacuum. Figure 8 illustrates this separation. Two black lines show the reflected solar illumination at an I/F of 0.25, and the emitted energy of a black body at 400 K, respectively. The red line illustrates the transmittance of the atmosphere for 10,000 ppm m of water vapor (approximately, ambient atmospheric water

vapor at sea level viewed through a 1 m path length). In order to prevent unknown atmospheric water vapor attenuation from influencing the laboratory calibration from 2,500 to 3,000 nm, we calibrate this region with the optical path in vacuum. We use both vacuum and free-air sources in the 2,000–2,500 nm range, providing an interval of overlap for assessing the consistency of the two calibration methods.

For wavelengths from 600 to 2,500 nm, the radiometry is traced to a lamp of known intensity calibrated by the National Institute of Standards and Technology (NIST). This lamp illuminates a reference panel of known reflectance set at a specific distance and angle, allowing the panel-leaving radiance to be calculated. Before and after each calibration measurement we obtain a dark data set with the chamber window port covered by an opaque cap. We then scan the instrument by rotating the optical bench assembly (OBA) within the chamber on a stage, so that the field of view sweeps across the illuminated panel. We extract the highest DN value for each cross-track position by fitting a curve to the rise and fall of the signal in the temporal dimension as the source sweeps over that field position. Independent measurements of the spectral reflectance and transmittance of the ground system elements (e.g., vacuum chamber window) in the HVM³ optical path are used to calculate the relationship between HVM³ DNs and the input radiance from the panel. This measurement is repeated multiple times to assess consistency, and duplicated for each of the two integration times. The same procedure is commonly used for terrestrial solar reflectance spectrometers, such as the airborne imaging spectrometer AVIRIS-NG (Chapman et al., 2019) or the orbital imaging spectrometer EMIT (Thompson et al., 2024). The 600–2,500 nm interval spans two regions where H₂O vapor attenuates the spectrum slightly: one at 1,380 nm and another at 1,880 nm. This results in minor divots that are visually apparent in the otherwise smooth RCC curves. We interpolate over these artifacts with local polynomials to obtain valid RCCs in these areas.

The midwave calibration is more complex; it requires an addition to the vacuum chamber described in detail by (McKinley, Hummel, et al., 2024). Briefly, the midwave source is a temperature-controlled blackbody target fabricated from black silicon, a material with a spectrally-flat emissivity above 99%. The source is placed within an external vacuum-sealed fixture that is bolted over the chamber port after removing the sapphire window, giving the instrument an unobstructed view of the source with the entire optical path in vacuum (Figure 9). This fixture contains a cold sleeve that minimizes any thermal emission except the black body itself. The black body is fitted with multiple contact temperature sensors to control the temperature and monitor for thermal gradients. Figure 10 shows the relationship between black body temperature and the radiance at the sensor. Red lines are model predictions based on Planck's law. Blue lines on the image indicate typical radiances expected for directly-illuminated and PSR targets. Due to the steepness of the black body radiance spectrum, we require multiple temperatures to span the whole interval. We measure the HVM³ radiometric response at temperatures from 250 to 400 K in multiple steps. Each step begins by raising the




Figure 9. Left: Overhead view of HVM³ instrument view of the black body target for midwave IR calibration. Right: Calibration chamber with black body assembly attached.

target to its set point temperature and waiting for the surface to equilibrate with no temperature gradients. We then acquire radiometric calibration measurements at both integration times. After these measurements, we quickly cool the source with liquid nitrogen through a heat exchanger to below 140 K, providing a cold surface for dark measurements. The calibration target does not fill the entire HVM³ field of view, so we sweep the field of view across the target, as for the NIST lamp, and then extract the brightest position using the same procedure. The result is the RCC and the flat field at all field points.

After assessing the data, we find that the collections with the black body at 350 K are most reliable. The spectrometer temperature during these acquisitions is closest to its optimal operating temperature of 170 K, minimizing uncertainty due to spectrometer thermal self emission. Moreover, two consecutive acquisitions at this temperature - one sweeping left to right, the other opposite - agree to within just 0.2% mean absolute error, indicating good temperature stability for the target. Consequently, we use the 350K data to form the RCC solution for HVM³, and reserve the other setpoints for validation. The black lines in Figure 10 represent measured radiances at 44 ms integrations based on the calibration result. The HVM³ data generally match Planck law predictions across the two and a half orders of magnitude in measured radiance. Small departures between red and black lines for certain temperatures are not systematically related to the radiance magnitude, as they would be the case of detector nonlinearity, or to wavelength, as they would be in the case of an RCC estimation error. A more likely cause is uncertainty in spectrometer self-emission and black body temperature for the lower-quality acquisitions.

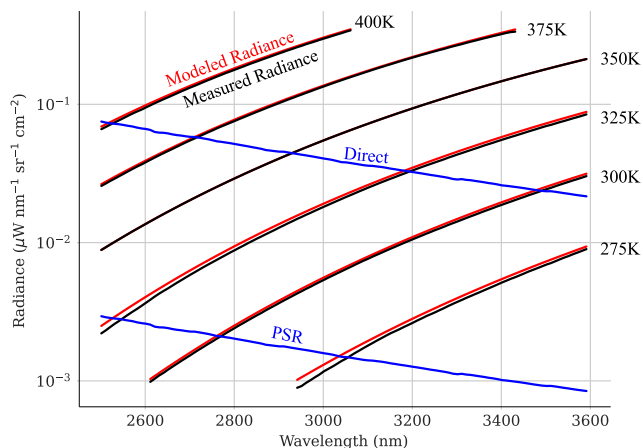


Figure 10. Black body emitted radiance predictions at different temperatures (red), compared with typical solar-reflected HVM³ targets (blue). Black lines correspond to HVM³ laboratory data. Blue lines indicate the reference radiance levels for direct solar and terrain-scattered (Permanently Shadowed Region) cases.

Figure 11 plots derived RCC vectors for midwave and shortwave calibrations on the same axis. The top curve corresponds to the RCCs for the short integration time. The two curves are nearly but not exactly proportional due to non-reciprocal electronic effects. The variable grating optical efficiency at different wavelengths causes a sinusoidal undulation in the RCCs. Two spikes at 1,190 nm and 2,170 nm are an expected loss in sensitivity near OSF seams. A minor bump in the RCC curves at 2,800 nm is likely a feature of the OSF transmissivity. Contamination was ruled out because this feature was perfectly stable across thermal cycles and decontamination heating periods. As further confirmation, a real 1% ice absorption feature appeared in testing at 3,100 nm after an anomaly in the environmental thermal control. We removed the errant absorption with standard decontamination procedures, through which the 2,800 feature remained unchanged. This is consistent with its interpretation as an intrinsic property of the instrument. For both integration times, the shortwave and midwave calibration measurements overlay within the overlap region from 2000 to 2,500 nm, indicating good agreement between the two solutions. It is notable that the red and black curves were derived using two completely different measurements with independent

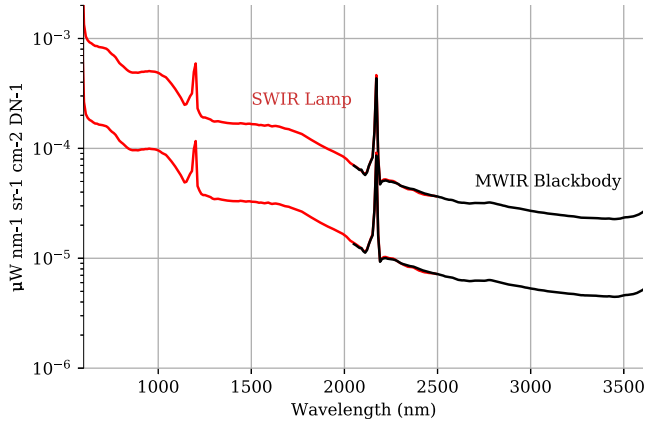


Figure 11. HVM³ calibration result showing Radiometric Calibration Coefficients (RCCs) inferred by two different laboratory methods: a calibrated incandescent source in open air (red line) and a temperature-controlled black body in vacuum (black line). The two offset lines represent RCCs for HVM³'s two different integration times. Spikes at 1,190 nm and 2,170 nm represent order sorting filter seams.

traceability paths - one from the NIST-calibrated illumination source, and the other from the calibration of temperature sensors on the blackbody target. Their agreement provides confidence in the accuracy of HVM³'s radiometric calibration.

4.2. Radiometric Sensitivity

We predict HVM³ radiometric sensitivity using a radiometric model describing the efficiency of each optical component, similar to previous missions like EMIT (Thompson et al., 2024). The total noise at each spatial pixel is a function $\sigma_{tot}(r)$ of the spectral channel given by detector row r :

$$\sigma_{tot}(r) = \sqrt{\sigma_p(r)^2 + \sigma_d^2 + \sigma_q^2 + \sigma_r^2 + \sigma_e^2} \quad \text{for} \quad (3)$$

$$\sigma_p(r) = \sqrt{s_{obs}(r) + s_{spc}(r)}$$

where photon shot noise $\sigma_p(r)$ is related to the observed scene signal s_{obs} in electrons, and the thermally-emitted signal s_{spc} from the spectrometer itself.

The symbol σ_d is dark current noise, σ_q is quantization noise, σ_r is read noise and σ_e is electronic noise. For observations of radiance $L_{obs}(r)$ in channel r , the corresponding signal in electrons at the detector during each integration time is:

$$s_{obs}(r) = kt_{int}a_{\omega}w_{obs}(r)L_{obs}(r)q_e(r)\Delta_{chn} \quad (4)$$

here k is a conversion factor defined as $k = 10^6 E^{-1}$, with photon energy E of hc/λ Joules, where h is Planck's constant and c is the speed of light in a vacuum. The symbols t_{int} and Δ_{chn} represent the integration time in seconds and the effective channel width in nm, respectively. The function $w_{obs}(r)$ represents the total optical throughput, including the losses due to the grating, OSF filters and mirrors, in channel r . The detector quantum efficiency $q_e(r)$ is also a function of the spectral channel. The symbol a_{ω} is the etendue of the scene photons. It is related to the detector area a_{det} , and the instrument F-number:

$$a_{\omega} = \frac{1}{4}a_{det}\pi F^{-2} \quad (5)$$

Emitted signal from the spectrometer interior follows a similar pattern, but because it is not dispersed, we integrate over the entire bandpass of the detector:

$$s_{spc}(r) = kt_{int}a_{\nu}w_{spc}(r)\int_{600nm}^{3600nm} B(T,\lambda)q_e(\lambda)d\lambda \quad (6)$$

where $B(T,\lambda)$ is the black body radiance evaluated at a specific wavelength λ . Here $w_{spc}(r)$ is the transmittance of the optical path for thermally-emitted photons from the spectrometer interior, and a_{ν} is the etendue:

$$a_{\nu} = \frac{1}{4}a_{det}\pi \quad (7)$$

This implicitly assumes a hemispherical view of a black body spectrometer interior, which is an approximation because much of the detector's local view is comprised of inefficient emitters. However, such differences are mathematically indistinguishable from changes to $w_{spc}(r)$, so for simplicity we bookkeep all differences there.

To determine HVM³ noise properties, we initialize the radiometric model from vendor and laboratory measurements of the efficiencies of different optical components. We then update w_{obs} and w_{spc} to the as-built condition using scene and thermal signals measured in the laboratory. The as-built dark noise are determined

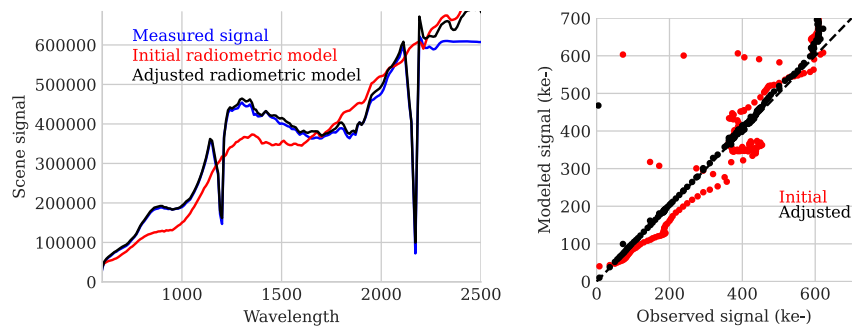


Figure 12. Validation of adjusted radiometric model. The model was updated using an independent observation acquired with a different integration time. Left: Adjustments improve the match between model predictions and measured signals. Right: The modified radiometric model fails only when the detector approaches saturation at 600 ke-.

from darks measured during the radiometric calibration process. We estimate the as-built throughput w_{obs} by comparing the measured signal from known radiometric sources with the predictions of the baseline radiometric model. The passively-cooled spectrometer will change temperature over the course of the mission. At temperatures higher than approximately 170 K, a non-negligible spectrometer self-emission signal radiates from the interior of the OBA directly on to the FPA. This self-emission has a distinct spectral character due to the attenuation by different OSF zones. We determine the as-built emissive throughput w_{spc} by assessing thermal emission using a dark scene and a range of different spectrometer temperatures. These noise and throughput measurements are incorporated into the HVM³ radiometric model. Figure 12 validates our radiometric model updates using a separate shortwave infrared calibration sequence. The validation test uses a separate reflectance panel and integration time from the data set we used to update the model. The updated model produces highly accurate predictions of the measured signal. As expected, it does not apply at signal levels near 600 ke- as the detector begins to saturate.

5. Lunar Trailblazer In-Flight Calibration Activities

Calibration is a process, not a result, as instrument properties can measurably over time. Spectral response can shift as components move under launch forces or gravity unloading. Radiometric response can change over hour-to-month timescales with drift in the thermal state of the detector and its electronics (Thompson et al., 2024). Contamination can reduce radiometric sensitivity. Thus, HVM³ will acquire calibration data in flight to monitor, and if necessary update, its laboratory calibration files. It will have many opportunities for this throughout the mission. After launch, Lunar Trailblazer will enter a multi-month cruise phase with month-long decontamination cycles to drive off any volatile deposits that have accumulated on HVM³. We will end each of these periods with calibration acquisitions to check for contamination and reassess instrument performance. Next, upon arrival at the Moon, Lunar Trailblazer will begin a period reduction phase to circularize its orbit. Here HVM³ will minimize calibration activities to free the spacecraft for regular propulsive maneuvers. After achieving science orbit, HVM³ will begin routine operations in which periodic calibration activities interspersed with lunar science data acquisitions. The following section describes each in-flight calibration activity in greater detail.

Updating the flat field is important because even minor, sub-percent deviations in the relative sensitivity of different elements can create along-track stripes in derived products. HVM³ has multiple strategies to measure the flat field in flight. First, long temporal sequences of science data can be used to average over terrain variability, creating a uniform stimulus for measuring the relative response of detector elements. In theory, with enough images, the randomly-arranged patches of light and shadow on the lunar terrain average out to a uniform gray. In practice, however, it can be challenging to obtain uniformity within a fraction of a percent for an accurate flat field. Such uniformity could require hundreds or thousands of images captured in a timespan where radiometric drift is negligible, a volume that is infeasible for Lunar Trailblazer. Fortunately, it has been shown that preprocessing the data with edge removal and a high-pass filter (Thompson et al., 2024) enables accurate measurement of the high spatial frequency components of the flat field, that is the pixel-to-pixel variability in individual detector elements, with a much smaller data set. This approach is

practical for HVM³ operations and should mitigate the element-wise radiometric drift that is a main cause of pushbroom stripes.

A second strategy for measuring the flat field involves changing the spacecraft orientation so that the slit is parallel to the direction of motion along the lunar surface. All FPA elements sweep out the same set of locations on the lunar surface so that their relative response can be measured. This technique, known colloquially as the side slither (Pesta et al., 2014), enables measurement of high- and low-frequency spatial components in a single acquisition. The execution is slightly different in deep space and lunar orbit. In deep space, with the Moon subtending a number of pixels on the focal plane, the spacecraft slews so that the field of view sweeps laterally across the center of the lunar disk, with the slit's linear field of view oriented parallel to the direction of motion. In lunar science orbit, the spacecraft simply yaws 90° for the duration of the acquisition. Note that, while the side slither has been used effectively in pushbroom instruments for decades, its application to HVM³'s instrument geometry is experimental. The center and periphery of the array cannot see exactly the same locations on the lunar surface due to the natural curvature of the slit's projection. Thus, the side slither also relies on some degree of spatial uniformity and averaging to achieve an accurate flat field. Thus, we will use this data judiciously.

In addition to the flat field, it is important to assess and update RCCs. HVM³ will acquire periodic views of reference stars to monitor for changes in sensitivity. Specifically, we will sweep the HVM³ field of view across a bright star with a known spectral flux, comparing predicted and measured irradiance on the instrument aperture. α Lyrae (Vega) and α Canis Majoris (Sirius) are good reference stars due to their high apparent magnitude and extensive prior study as radiometric standards (Cohen et al., 1992). Our calculations suggest that both should be visible to a slow pushbroom sweep of the target using HVM³'s long integration time. This is particularly critical during the initial post-launch phase of the mission, where we monitor for the presence of volatiles that have condensed on the instrument after delivery. Ratios of star irradiance spectra acquired after month-long decontamination cycles will be used to monitor for any change in the water absorption bands.

We will monitor spectral calibration in flight by fitting the positions of known spectral features. The Earth itself has sharp atmospheric features due to gas absorption by water vapor, O₂, CO₂, and other gases. At regular intervals during cruise and science operations phases, we will sweep the Earth disk across at three points (left, center, and right) within the instrument field of view, and fit the wavelength calibration using the known spectral positions of these atmospheric features. This procedure, described in Thompson et al. (2024) for the EMIT mission, can optimize both wavelength center position and the SRFs. Another useful check on instrument calibration is a minor absorption feature in the solar irradiance curve near 800 nm. This feature will appear in all HVM³ science data, providing a convenient standard for wavelength calibration with every acquisition.

The spatial response functions are unlikely to change measurably over the course of the mission. Nevertheless, we will validate instrument spatial response properties using views of the lunar horizon. The horizon provides a crisp edge that can be used to confirm both peak and tail properties of the system spatial response functions.

6. Performance Model

The as-built radiometric model predicts HVM³'s SNR for typical lunar surfaces. Exploiting Kirchhoff's law, we use a simple model for the radiance vector at sensor, \mathbf{L}_{obs} :

$$\mathbf{L}_{\text{obs}} = \mathbf{F}\pi^{-1} \cos(\theta) \circ \boldsymbol{\rho} + \mathcal{B}(T) \circ (1 - \rho) \quad (8)$$

here we use bold symbols to signify vector-valued spectral quantities. Adjacency of bold symbols indicates a vector or matrix product while the \circ operator signifies element-wise multiplication. The symbol \mathbf{F} is the annual average solar irradiance at 1 AU, and θ is the local solar zenith angle, the angle of incident light. The symbol $\boldsymbol{\rho}$ is the hemispherical directional surface reflectance (Schaepman-Strub et al., 2006), for which we use a RELAB database reference measurement, BKR1LR117, an Apollo lunar regolith sample (Pieters, 1983). We use $\mathcal{B}(T)$ to represent the black body radiance at temperature T . Figure 13 shows the resulting SNR for observations at a spectrometer temperature of 186 K. The direct solar case uses a solar zenith angle of 85°, a local surface temperature of 300 K, and an integration time of 45 ms. This is near-worst case as more typical smaller zenith angles will have greater signal. The PSR case uses a local surface temperature of 80 K and a longer 227 ms integration time.

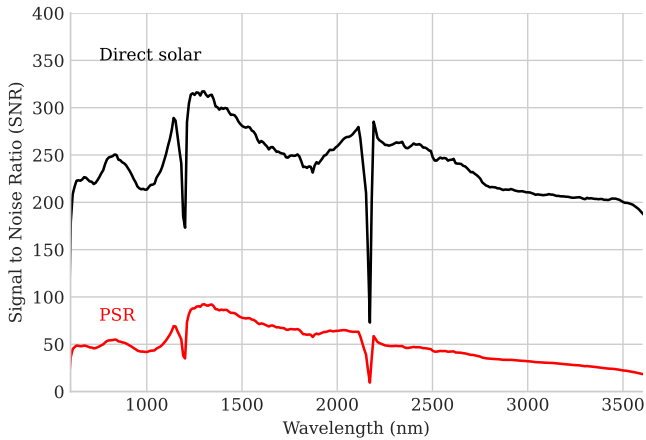


Figure 13. HVM³ signal to noise ratio. Lines indicate the relationship for direct solar observations at a solar zenith angle of 85° and a Permanently Shadowed Region observation with an irradiance of 1 W m⁻².

Figure 14 shows a graph of noise sources, including scene shot noise and spectrometer emission for an ideal spectrometer temperature of 170 K. We predict this temperature will be achieved during orbits with a high solar beta angle. We assume a direct solar observation with a local solar zenith angle of 85°. This system is shot noise limited; read noise is a tiny fraction of scene noise, and the noise from spectrometer thermal emission is even smaller. The situation is different for the PSR observation in the right subfigure, which has a much dimmer terrain-reflected irradiance of 1 W m⁻² and a longer integration time of 227 ms. The longer integration time increases the spectrometer thermal noise and the dark current noise. Simultaneously, the signal from the scene is reduced. As a result, the read noise, spectrometer thermal emission, and scene shot noise are all similar in magnitude. This is the case for ideal operating conditions with a spectrometer temperature of 170 K; if it rose above this level, the emitted signal would fast become the dominant noise source.

After determining the instrument's radiometric sensitivity, we assess HVM³ performance against the original science objectives. We define three reference cases representing different kinds of HVM³ observations. Each corre-

sponds to a hypothesis test using the HVM³ measurement to address a specific mission science question. These reference observations are stressing cases meant to simulate the most challenging discriminations we expect HVM³ to make. They all involve the signature of lunar water volatiles expressed solely in the 2,500–3,600 nm range. Naturally, the HVM³ instrument is also capable of measuring volatile and mineral signatures at shorter wavelengths. However, the 2,500–3,600 nm region is especially challenging and important for our investigation, so we focus our assessment on this wavelength range.

Following from our science objectives, the three reference cases are: (a) the discrimination of ice and molecular H₂O in PSRs, via scattered light; (b) the characterization of ice and molecular H₂O in midlatitude and highland areas illuminated at 85° solar zenith; and (c) detection of band depth change over time in adsorbed molecular H₂O in midlatitude areas illuminated at 85° solar zenith. The band depth for all features and changes is 1%; that is, we aim to discriminate ice and adsorbed molecular H₂O in PSRs for a 1% absorption depth. This corresponds roughly to ~100 ppm and ~1,000 ppm for sunlit and shadowed cases, respectively. This criterion is stricter than the >0.5% mass fraction of ice deemed useful for in situ resource utilization (Hayne et al., 2014), and provides a comfortable floor for detecting trace water elsewhere.

Table 2 shows the key parameters for each of the three reference cases. We define the null hypothesis to be the radiance from the BKR1LR117 spectrum attenuated by an H₂O absorption with 1% band depth. The alternative hypothesis for midlatitude and PSR cases is that the spectral absorption feature is water ice with 1% band depth. For the change detection task, the alternative hypothesis is that the adsorbed molecular water absorption grows in size to a 2% feature. For the solar-illuminated cases, we use a solar zenith angle of 85°. Our PSR illumination estimate




Figure 14. Left: HVM³ noise sources for direct solar observation at optimal spectrometer temperature. Right: HVM³ noise sources for Permanently Shadowed Region observations at optimal spectrometer temperature. Shot noise from scene photons appears in blue, while spectrometer-emitted photon shot noise appears in red.

Table 2
HVM³ Reference Scenarios

Scenario	H_0	H_1	Irradiance	Binning	Integration	p value
I. PSR	1% H ₂ O	1% Ice	Scattered	45X	227 ms	0.019
II. Midlatitude	1% H ₂ O	1% Ice	Direct	9X	44 ms	<0.01
III. Change	1% H ₂ O	2% H ₂ O	Direct	9X	44 ms	<0.01

Note. Direct illumination assumes the solar zenith angle is 85°. Scattered illumination assumes terrain-reflected irradiance of $1\mu\text{Wcm}^{-2}\text{nm}^{-1}\text{sr}^{-1}$.

assumes a 1 W m^{-2} incident irradiance, which is consistent with measurements by the LROC instrument as well as modeling activities in recent literature (Mazarico et al., 2011). For simplicity, we assume lunar thermal emissivity effects are negligible at these latitudes. Each scenario permits some level of spatial binning to improve SNR. Pixels in the midlatitude observations are binned spatially into 3×3 blocks, providing a 240×240 m footprint with a factor of three noise improvement relative to the single-pixel spectrum. PSR observations use even coarser aggregation; they begin with a $5\times$ longer integration time, and are binned $15\times$ in the cross-track direction and $3\times$ in the downtrack direction to yield an effective observation footprint of approximately 1200×1200 m.

To assess the sensitivity of HVM³ measurements to lunar volatiles, we must first specify a retrieval algorithm. We exploit the fact that for small optical path lengths, the Beer-Lambert absorption model can be approximated accurately using a first-order Taylor series expansion. For an optical path length ℓ close to zero, H₂O absorption coefficient κ , absorption modifies the background upwelling radiance \mathbf{L}_b to the observed radiance \mathbf{L}_{obs} according to:

$$\mathbf{L}_{obs} = \mathbf{L}_b e^{-\ell\kappa} \approx \mathbf{L}_b - \mathbf{L}_b \ell \kappa \quad (9)$$

This is highly accurate for low temperatures without surface-emitted photons. A similar linearization would apply for higher temperatures where changes in surface emissivity could also play a role. However, since our most stressing observation cases involve cold scenes, we can safely ignore surface emission. For a radiance measurement vector \mathbf{x} perturbed by random instrument noise ϵ , the test to discriminate between the null hypothesis \mathbf{H}_0 and the test hypothesis \mathbf{H}_1 must consider the two alternatives:

$$\mathbf{H}_0 : \mathbf{x} = \mathbf{L}_0 + \epsilon \quad \mathbf{H}_1 : \mathbf{x} = \mathbf{L}_1 + \epsilon \quad (10)$$

The instrument noise ϵ is a vector-valued random variable with as many dimensions as instrument channels, distributed following a zero mean Gaussian distribution with covariance matrix Σ :

$$\epsilon \sim \mathbf{N}(\mathbf{0}, \Sigma) \quad (11)$$

We can calculate Σ for any radiance being observed using the HVM³ instrument model. The symbol \mathbf{L}_0 represents the null case radiance. Following a linearized Beer-Lambert relation of Equation 9, we write:

$$\mathbf{L}_0 = \mathbf{L}_b - \mathbf{L}_b \ell_0 \kappa_0 \quad (12)$$

with ℓ_0 representing the absorption path length for the null hypothesis, and κ_0 representing the vector of absorption coefficients for null hypothesis. Similarly, \mathbf{L}_1 is the test case radiance given by:

$$\mathbf{L}_1 = \mathbf{L}_b - \mathbf{L}_b \ell_1 \kappa_1 \quad (13)$$

with ℓ_1 and κ_1 representing the path length and vector of absorption coefficients respectively. Figures 15 and 16 show the hypothesis tests for species discrimination and band depth change measurement plotted as continuum-relative absorption features.

For convenience, we can rewrite Equation 10 in the following form:




Figure 15. Reference case for the discrimination of adsorbed molecular water and water ice. Left: The two hypotheses shown as continuum-relative absorption features. The black line indicates absorption by adsorbed molecular H_2O . The red line indicates absorption by water ice. Right: The β coefficients for the hypothesis test.

$$H_0 : \mathbf{x} = \mathbf{L}_0 + \epsilon \quad H_1 : \mathbf{x} = \mathbf{L}_0 + \mathbf{t} + \epsilon \quad (14)$$

where \mathbf{t} is a target signature given by:

$$\mathbf{t} = \mathbf{L}_b \ell_0 \kappa_0 - \mathbf{L}_b \ell_1 \kappa_1 \quad (15)$$

It is a well-established result (Manolakis et al., 2001) that the most powerful detection rule for the system in Equation 14 is a likelihood ratio test called a matched filter given by the expression:

$$\alpha(\mathbf{x}) = \frac{(\mathbf{x} - \mathbf{L}_0)^T \Sigma^{-1} \mathbf{t}}{\mathbf{t}^T \Sigma^{-1} \mathbf{t}} = (\mathbf{x} - \mathbf{L}_0)^T \beta \quad (16)$$

The retrieval algorithm is thus a linear operator β applied to $(\mathbf{x} - \mathbf{L}_0)$ that returns a scalar value. The expression for β is normalized to return a value of unity when \mathbf{x} is exactly equal to \mathbf{L}_1 . Note that this algorithm is not intended to be a general purpose method for retrieving water or ice on the lunar surface. The estimator is optimal only for spectra where the first order Taylor expansion of Equation 9 is valid, in other words, for spectra that are similar to \mathbf{L}_0 . This is certainly the case for the noisy null and test case spectra defined above.

For a classical hypothesis test, a statistically confident detection holds the risk of a false positive below $p = 0.05$, that is a five percent probability. In other words, the noise must be low enough that the matched filter, applied to a noisy observation of \mathbf{L}_0 , yields a value greater than unity no more than 5% of the time. Combining Equations 10 and 16, we have:

$$\alpha(\mathbf{L}_0 + \epsilon) = (\mathbf{L}_0 + \epsilon - \mathbf{L}_0)^T \beta = \epsilon \beta \quad (17)$$




Figure 16. Reference case for the measurement of changes in band depth over time. Left: The two hypotheses shown as continuum-relative absorption features. Right: The β coefficients for the hypothesis test.

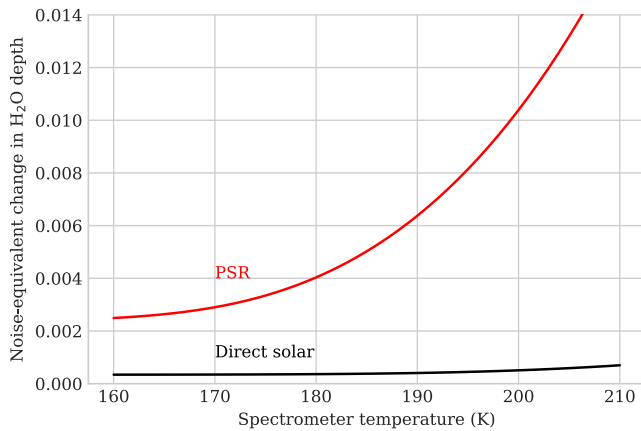


Figure 17. HVM³ detection sensitivity as a function of spectrometer temperature. The vertical axis shows the one-sigma change in measured band depth induced by instrument noise. HVM³ achieves optimal sensitivity when the spectrometer temperature is at or below 170 K.

Using a linear transformation of the Gaussian random variable ϵ :

$$\alpha(\mathbf{L}_0 + \epsilon) \sim \mathbf{N}(\mathbf{0}, \beta^T \Sigma \beta) \quad (18)$$

This expression indicates the noise-induced variance in the retrieval result, from which the probability of exceeding unity can be calculated using the normal conditional distribution function. The associated p -values appear as the final columns in Table 2. The p -values are smaller than 0.05 in all cases; by this conventional standard, the detections are statistically confident.

The instrument provides statistically-confident detections in all three cases for the nominal spectrometer temperature of 186 K. As the spectrometer temperature varies over the mission due to changes in solar beta angle, the differing thermal photon shot noise will change the sensitivity. To assess this relationship we calculated the noise-equivalent change in the depth of an H₂O absorption feature as a function of spectrometer temperature, using a linear detection rule. Figure 17 shows the result. The directly-illuminated observation in black is insensitive to thermal emission. However, PSR performance, with its longer integration time and dimmer scene, degrades

significantly for spectrometer temperatures above 190 K. This is consistent with the findings of the instrument noise model. Consequently, spectrometer temperature will be considered in spacecraft planning to optimize the acquisition of data during favorable beta angles. Fortunately, these changes in temperature should not affect the spectral calibration. Laboratory measurements indicate that the all-aluminum spectrometer retains its alignment across such changes in temperature.

The statistical confidences reported in Table 2 indicate the radiometric sensitivity of the HVM³ instrument design, but they may underestimate some practical challenges of measuring PSRs. Absorptions are only apparent in the ratio of upwelling to downwelling energy. Consequently, in addition to measuring the signal from the PSR itself, it will also be necessary to measure the radiance of the reflecting surface that is illuminating the PSR. Typically, these are crater walls in direct sunlight that may contain their own volatiles. Other uncertainties in data interpretation include the nature of volatile absorption coefficients, which can change based on temperature and intimate mixing effects. Radiometric and spectral calibration could also be challenging. There are no spatially extended radiometric standards visible on the lunar surface, but stellar observations are planned to track relative radiometry and confirm successful post-launch decontamination of optical components. The Earth itself provides an excellent spectral calibration source due to the presence of known atmospheric absorption features (Thompson et al., 2024). A solar feature near 800 nm provides another absolute reference. These and other in-flight calibration activities, described above in Section 5, will continuously validate HVM³ performance.

7. Conclusions

Drawing from previous spectrometers like M³ and UCIS, HVM³ was designed with specific objectives related to measurement of subtle volatile absorption features on the lunar surface. Laboratory data combined with performance models indicate that HVM³ as-built has the radiometric precision to achieve this task. HVM³ exhibits high spatial uniformity, high radiometric sensitivity, and a calibration result that agrees across two disparate traceability paths. Thus, HVM³ as-built can measure small perturbations of volatile absorption features with high confidence under direct solar illumination. A special long integration mode, combined with spatial coaddition, enables statistically confident measurements at the 1% level in PSRs with terrain-scattered illumination of 1 W m⁻². Water ice detections at even lower illuminations are accessible by combining information from the 1,500 nm and 2,000 nm water absorptions (high signal spectral regions) or co-adding multiple images acquired at different times.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

The HVM³ calibration code (Thompson, 2023) and preliminary calibration data are available at <https://github.com/LunarTrailblazr/HVM3-Calibration> with a permanent archival copy at <https://doi.org/10.5281/zenodo.14963157>. Data is available from (Thompson, 2025).

Acknowledgments

We acknowledge the support of the NASA Small Innovative Missions for Planetary Exploration (SIMPLEX) program. A portion of the research was carried out at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration (80NM0018D0004). Copyright 2025 California Institute of Technology. All Rights Reserved.

References

- Bandfield, J. L., Poston, M. J., Klima, R. L., & Edwards, C. S. (2018). Widespread distribution of oh/h₂o on the lunar surface inferred from spectral data. *Nature Geoscience*, *11*(3), 173–177. <https://doi.org/10.1038/s41561-018-0065-0>
- Bender, H. A., Smith, C. D., Ehlmann, B. L., Thompson, D. R., Vinckier, Q. P., & Mouroulis, P. (2022). Optical design and performance of the lunar trailblazer high-resolution volatiles and minerals moon mapper (hvm3). *Imaging spectrometry xxv: Applications, sensors, and processing*, *12235*, 11–17.
- Boardman, J. W., Pieters, C., Green, R., Lundeen, S., Varanasi, P., Nettles, J., et al. (2011). Measuring moonlight: An overview of the spatial properties, lunar coverage, selenolocation, and related level 1b products of the moon mineralogy mapper. *Journal of Geophysical Research*, *116*(E6), E00G14. <https://doi.org/10.1029/2010je003730>
- Bowles, N. E., Ehlmann, B. L., Evans, R., Warren, T., Eshbaugh, H., King, G., et al. (2025). The lunar trailblazer lunar thermal mapper instrument. *Journal of Geophysical Research: Planets*. in review. <https://doi.org/10.1029/2025JE009333>
- Candès, E. J., Li, X., Ma, Y., & Wright, J. (2011). Robust principal component analysis? *Journal of the ACM*, *58*(3), 1–37. <https://doi.org/10.1145/1970392.1970395>
- Chapman, J. W., Thompson, D. R., Helmlinger, M. C., Bue, B. D., Green, R. O., Eastwood, M. L., et al. (2019). Spectral and radiometric calibration of the next generation airborne visible infrared spectrometer (Avisir-ng). *Remote Sensing*, *11*(18), 2129. <https://doi.org/10.3390/rs11182129>
- Chauhan, P., Chauhan, M., Verma, P. A., Sharma, S., Bhattacharya, S., Dagar, A. K., et al. (2021). Unambiguous detection of oh and h₂o on the Moon from Chandrayaan-2 imaging infrared spectrometer reflectance data using 3 μm hydration feature. *Current Science*, *121*(3), 391–401. <https://doi.org/10.18520/cs/v121/i3/391-401>
- Clark, R. N. (2009). Detection of adsorbed water and hydroxyl on the Moon. *Science*, *326*(5952), 562–564. <https://doi.org/10.1126/science.1178105>
- Colaprete, A., Schultz, P., Heldmann, J., Wooden, D., Shirley, M., Ennico, K., et al. (2010). Detection of water in the LCROSS ejecta plume. *Science*, *330*(6003), 463–468. <https://doi.org/10.1126/science.1186986>
- Cohen, M., Walker, R. G., Barlow, M. J., & Deacon, J. R. (1992). Spectral irradiance calibration in the infrared. 1. Ground-based and iras broad-band calibrations. *The Astronomical Journal*, *104*(4), 1650–1657. <https://doi.org/10.1086/116349>
- Dapremont, A. M., Klima, R. L., Wilk, K., Ehlmann, B. L., Edwards, C. S., Hanna, K. L. D., et al. (2025). Visible-shortwave infrared (VSWIR) spectral parameters for the lunar trailblazer high-resolution volatiles and minerals Moon Mapper (HVM3). *Earth and Space Science*. in review. <https://doi.org/10.1029/2025EA004557>
- Ehlmann, B. L., Klima, R. L., Seybold, C., Klesh, A., Bennett, C. L., Bowles, N., et al. (2025). The lunar trailblazer mission: Science motivation and implementation of a pioneering small satellite for lunar water and lunar geology in the NASA SIMPLEX program. *Journal of Geophysical Research: Planets*. in review. <https://doi.org/10.1029/2025JE009300>
- Hayne, P. O., Ingersoll, A. P., & Paige, D. A. (2014). New approaches to lunar ice detection and mapping.
- Honniball, C. I., Lucey, P. G., Li, S., Shenoy, S., Orlando, T. M., Hibbitts, C. A., et al. (2021). Molecular water detected on the sunlit Moon by SOFIA. *Nature Astronomy*, *5*(2), 121–127. <https://doi.org/10.1038/s41550-020-01222-x>
- Laferriere, K., Sunshine, J., & Feaga, L. (2022). Variability of hydration across the southern hemisphere of the Moon as observed by deep impact. *Journal of Geophysical Research: Planets*, *127*(8), e2022JE007361. <https://doi.org/10.1029/2022je007361>
- Li, S., Lucey, P. G., Milliken, R. E., Hayne, P. O., Fisher, E., Williams, J.-P., et al. (2018). Direct evidence of surface exposed water ice in the lunar polar regions. *Proceedings of the National Academy of Sciences*, *115*(36), 8907–8912. <https://doi.org/10.1073/pnas.1802345115>
- Li, S., & Milliken, R. E. (2017). Water on the surface of the moon as seen by the moon mineralogy mapper: Distribution, abundance, and origins. *Science Advances*, *3*(9), e1701471. <https://doi.org/10.1126/sciadv.1701471>
- Manolakis, D. G., Siracusa, C., Marden, D., & Shaw, G. A. (2001). Hyperspectral adaptive matched-filter detectors: Practical performance comparison. *Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII*, *4381*, 18–33. <https://doi.org/10.1117/12.437006>
- Mazarico, E., Neumann, G., Smith, D., Zuber, M., & Torrence, M. (2011). Illumination conditions of the lunar polar regions using lola topography. *Icarus*, *211*(2), 1066–1081. <https://doi.org/10.1016/j.icarus.2010.10.030>
- McKinley, I. M., Hummel, C. D., Keller, J., Moore, B. D., & Rodriguez, J. I. (2024). Thermal vacuum testing of the hvm3 instrument and ucis-moon imaging spectrometer (pp. 1–10).
- McKinley, I. M., Rodriguez, J. I., Allen, J. D., Keller, J., Mok, M. A., & Chan, L. (2024). Thermal design of the high-resolution volatiles and minerals Moon Mapper (hvm3) instrument on the lunar trailblazer mission (ltb). *IOP Conference Series: Materials Science and Engineering*, *1301*(1), 012012. <https://doi.org/10.1088/1757-899x/1301/1/012012>
- Pesta, F., Bhatta, S., Helder, D., & Mishra, N. (2014). Radiometric non-uniformity characterization and correction of landsat 8 oli using Earth imagery-based techniques. *Remote Sensing*, *7*(1), 430–446. <https://doi.org/10.3390/rs70100430>
- Pieters, C. M. (1983). Strength of mineral absorption features in the transmitted component of near-infrared reflected light: First results from relab. *Journal of Geophysical Research*, *88*(B11), 9534–9544. <https://doi.org/10.1029/jb088ib11p09534>
- Pieters, C. M., Goswami, J., Clark, R., Annadurai, M., Boardman, J., Buratti, B., et al. (2009). Character and spatial distribution of OH/H₂O on the surface of the Moon seen by M³ on Chandrayaan-1. *Science*, *326*(5952), 568–572. <https://doi.org/10.1126/science.1178658>
- Schaepman-Strub, G., Schaepman, M. E., Painter, T. H., Dangel, S., & Martonchik, J. V. (2006). Reflectance quantities in optical remote sensing—definitions and case studies. *Remote Sensing of Environment*, *103*(1), 27–42. <https://doi.org/10.1016/j.rse.2006.03.002>
- Sunshine, J. M., Farnham, T. L., Feaga, L. M., Groussin, O., Merlin, F., Milliken, R. E., & A’Hearn, M. F. (2009). Temporal and spatial variability of lunar hydration as observed by the deep impact spacecraft. *Science*, *326*(5952), 565–568. <https://doi.org/10.1126/science.1179788>
- Thompson, D. R. (2023). HVM³ code repository [Software]. <https://doi.org/10.5281/zenodo.14963157>
- Thompson, D. R. (2025). HVM³ calibration code and data release 1.0 [Dataset]. *Zenodo*. <https://doi.org/10.5281/zenodo.14963157>

- Thompson, D. R., Green, R. O., Bradley, C., Brodrick, P. G., Mahowald, N., Dor, E. B., et al. (2024). On-orbit calibration and performance of the emit imaging spectrometer. *Remote Sensing of Environment*, *303*, 113986. <https://doi.org/10.1016/j.rse.2023.113986>
- Van Gorp, B., Mouroulis, P., Blaney, D., Green, R. O., Ehlmann, B. L., & Rodriguez, J. I. (2014). Ultra-compact imaging spectrometer for remote, in situ, and microscopic planetary mineralogy. *Journal of Applied Remote Sensing*, *8*(1), 084988. <https://doi.org/10.1117/1.jrs.8.084988>
- Wöhler, C., Grumpe, A., Berezhnoy, A. A., & Shevchenko, V. V. (2017). Time-of-day-dependent global distribution of lunar surficial water/hydroxyl. *Science Advances*, *3*(9), e1701286. <https://doi.org/10.1126/sciadv.1701286>